

Appl. No. 09/401,616  
Amdt. dated January 21, 2005  
Reply to Office action of October 21, 2004

**Amendments to the Claims:**

This listing of claims will replace all prior versions, and listings, of claims in the application:

**Listing of Claims:**

1. (Currently amended) A method for value sampling for monitoring the performance of a program being executed on a computer system, comprising the steps of:
  - executing the program on a computer system, the program having object code instructions;
  - at time-based intervals interrupting execution of the program, including delivering a first interrupt; and
  - in response to at least a subset of the first interrupts, storing at least one data value of interest in a first database, the at least one data value of interest being associated with a particular object code instruction of the program, the particular object code instruction being executed by the computer when the interrupt occurs, wherein the particular object code instruction of the program remains unmodified.
2. (Currently amended) The method of claim 1 wherein the intervals are random time intervals.
3. (Currently amended) The method of claim 1 wherein the intervals have a constant time period.
4. (Original) The method of claim 1 further comprising the step of:
  - specifying an object code instruction of interest, and wherein the at least one data value of interest is associated with the specified object code instruction of interest.

Appl. No. 09/401,616  
Amdt. dated January 21, 2005  
Reply to Office action of October 21, 2004

5. (Original) The method of claim 1 wherein the data value of interest is an operand.
6. (Original) The method of claim 1 wherein the data value of interest is a result of the execution of the instruction.
7. (Original) The method of claim 1 wherein the particular object code instruction is associated with a memory address, the memory address being stored in a program counter, and wherein said step of storing stores the associated memory address and the at least one data value of interest in the database.
8. (Original) The method of claim 1 wherein the program includes object code instructions corresponding to a shared library, and said step of storing stores at least one data value of interest associated with at least one object code instruction derived from the shared library.
9. (Original) The method of claim 1 wherein the program includes object code instructions corresponding to at least one kernel instruction, and said step of storing stores at least one data value of interest associated with at least one kernel instruction.
10. (Original) The method of claim 1 further comprising the step of:  
specifying a set of object code instructions of interest, and  
wherein said step of storing is performed only when the program is interrupted at a point associated with at least one of the set of object code instructions of interest.
11. (Original) The method of claim 1 wherein said step of storing stores a current interrupt level.

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

12. (Original) The method of claim 1 wherein said step of storing stores physical addresses for load and store instructions to provide information about the number and sources of memory references.

13. (Original) The method of claim 1 wherein said step of storing stores at least one data value that identifies the destination of at least one instruction that transfers control flow as the at least one data value of interest.

14. (Previously presented) The method of claim 13 wherein said at least one instruction that transfers control flow includes any of a conditional branch instruction, an unconditional branch instruction, a jump instruction, a call subroutine instruction, or a subroutine return instruction, and the destination is identified by an associated destination address.

15. (Original) The method of claim 13 wherein the at least one instruction is a conditional branch instruction, and the destination is identified by a bit.

16. (Original) The method of claim 1 wherein said step of storing includes the steps of:

identifying at least one issue block of instructions;

interpreting the instructions of the at least one issue block; and

storing at least one data value of interest associated with at least one interpreted instruction.

17. (Original) The method of claim 16 wherein said step of storing stores the at least one data value of interest after interpreting each instruction of said at least one issue block.

18. (Original) The method of claim 16 wherein said step of interpreting emulates a machine language instruction set of the computer system.

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

19. (Original) The method of claim 16 wherein the interpreter updates a state of the interrupted program as though each interpreted instruction had been directly executed by the computer system.

20. (Original) The method of claim 16 wherein the interpreter interprets both kernel and user code.

21. (Original) The method of claim 16 wherein the interpreter does not execute particular instructions.

22. (Original) The method of claim 1 wherein the particular object code instruction is an interrupted instruction, the particular object code instruction being associated with a memory address, and said step of storing includes the steps of:  
storing the memory address and the interrupted instruction;  
configuring a second interrupt to be delivered after a predetermined number of instructions are executed; and  
in response to the second interrupt,  
deactivating the second interrupt, and  
storing the memory address and the at least one data value of interest for the associated instruction in the first database.

23. (Original) The method of claim 22 wherein the predetermined number of instructions is equal to a number of instructions of an issue block.

24. (Original) The method of claim 1 wherein the particular object code instruction is an interrupted instruction, the particular object code instruction being associated with a memory address, and said step of storing includes the steps of:  
storing the memory address and a set of object code instructions including the interrupted instruction;  
configuring a second interrupt to be delivered after a predetermined number of instructions; and

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

in response to the second interrupt,  
deactivating the second interrupt, and  
storing the memory address and the at least one data value of  
interest for the interrupted instruction in the first database.

25. (Original) The method of claim 24 further comprising the step of:  
analyzing the interrupted instruction to determine the at least one data  
value of interest to store in the first database.

26. (Original) The method of claim 1 wherein the particular object code  
instruction is associated with a set of object code instructions, the set of object  
code instructions also including an interrupted instruction, the interrupted  
instruction also being associated with a memory address, and said step of storing  
includes the steps of:

storing the memory address and the set of object code instructions  
including the interrupted instruction;  
configuring a second interrupt to be delivered after a predetermined  
number of events; and  
in response to the second interrupt,  
deactivating the second interrupt, and  
storing the memory address and the at least one data value of  
interest for at least one instruction of the set of object code  
instructions in the first database.

27. (Original) The method of claim 26 wherein the set of object code  
instructions is an issue block.

28. (Original) The method of claim 26 wherein the predetermined number of  
events is a predetermined number of instruction executions.

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

29. (Original) The method of claim 26 wherein the predetermined number of events is a predetermined number of clock cycles.
30. (Original) The method of claim 1 further comprising the steps of:  
periodically storing the at least one data value of interest of the first database in a second database.
31. (Original) The method of claim 30 further comprising the step of:  
generating at least one value profile for the at least one data value of interest in the second database.
32. (Original) The method of claim 1 wherein the at least one data value of interest has a plurality of data values of interest, and said step of storing stores a tuple of the plurality of the data values of interest and a memory address in the first database.
33. (Original) The method of claim 32 wherein one of the data values of interest is a particular data value of interest and the tuple includes a count that is associated with the particular data value of interest, and  
said step of storing further includes the steps of:  
identifying the particular data value of interest in the first database;  
and  
incrementing the count associated with the particular data value of interest.
34. (Original) The method of claim 32 wherein the data values of interest stored in the tuple include a value stored in a register specified by the interrupted instruction.

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

35. (Original) The method of claim 32 wherein the data values of interest stored in the tuple include the value stored in a destination register specified by the instruction and at least one other value stored in another register.

36. (Original) The method of claim 1 wherein the particular object code instruction is part of an issue block of instructions, and said step of storing stores one or more data values of interest for one or more instructions of the issue block, the data values of interest including a value of a register specified by one of the instructions of the issue block, including a value stored in a return address register and a value stored in a memory location in a current stack frame.

37. (Original) The method of claim 32, wherein one of the data values of interest is a particular data value of interest and the tuple is a particular tuple having particular values of interest, the tuple also having a count that is associated with all the particular data values of interest of the particular tuple, and said step of storing further includes the steps of:

identifying the particular tuple with the particular data values of interest in the first database based on the at least one data value of interest;  
and

incrementing the count associated with the particular tuple when the at least one data value of interest is the same as the particular data values of interest of the particular tuple.

38. (Original) The method of claim 1 wherein said step of storing stores a program counter value which invoked a function containing the profiled instruction and a value stored in a destination register as correlated values.

39. (Original) The method of claim 1 wherein said step of storing stores a return address which invoked a function containing the profiled instruction and a value stored in a destination register as correlated values, and further comprising the step of:

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

determining if the correlated values are already stored in the first database, and if so incrementing a counter associated with the correlated values, otherwise storing the correlated values in the first database, whereby the correlated values represent a profile of information to call sites.

40. (Original) The method of 1 further comprising the step of:  
optimizing the program of object code instructions based on the at least one data value of interest.

41. (Original) The method of claim 1 wherein the at least one data value of interest is context information, and said step of optimizing optimizes the program of object code instructions based on the context information.

42. (Original) The method of claim 1 wherein said step of storing stores a set of particular data values of interest for the particular object code instruction, and further comprising the steps of:

generating another program of object code instructions from a set of object code instructions such that separate, specialized versions of object code are generated when at least one of the set of particular data values of interest has a predetermined amount of invariance.

43. (Previously presented) The method of claim 42 wherein the set of particular data values of interest include referenced addresses to a memory, and said step of generating another program of object code instructions generates object code instructions to prefetch instructions from at least one of the referenced addresses when the at least one of the referenced addresses has a sufficient invariance to decrease a latency of a memory reference.

44. (Original) The method of claim 42 wherein the set of particular data values of interest include an expected value of an input value for a particular instruction,



**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

and said step of generating generates another program of object code instructions to execute the particular instruction using the expected value, and that generates object code instructions that determine an actual value of the input value for the particular instruction, and compares the actual value to the expected value to determine whether the particular instruction was executed with an appropriate value.

45. (Original) The method of claim 1 wherein said step of storing includes the step of:

applying a predetermined function to the at least one data value of interest before the data value of interest is stored in the first database.

46. (Original) The method of claim 1 further comprising the step of:

identifying an instruction of interest having at least one operand; and said step of storing including the steps of:

generating at least one projected value from the at least one operand by applying a function to the at least one operand of the instruction of interest; and

storing the at least one projected value in the first database.

47. (Original) The method of claim 32, wherein one of the data values of interest is a particular data value of interest and the tuple is a particular tuple having particular values of interest, the tuple also storing a functional value, wherein the first database stores a set of tuples that are associated with different particular data values of interest associated with the particular object code instruction,

for the particular object code instruction, updating the functional value in the particular tuple based on the at least one data value of interest associated with the particular object code instruction and at least one data value of interest of each tuple in the set of tuples.

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

48. (Original) The method of claim 16 further including the steps of:  
receiving a downloaded script; and  
wherein said step of storing includes the steps of:  
after interpreting one of the instructions, executing the downloaded  
script to determine and store the at least one data value of  
interest.
49. (Original) The method of 16 further including the steps of:  
receiving an interpretable program having interpretable instructions;  
receiving a downloaded script that causes a user-mode trap to be sent to  
the interpretable program;  
executing the interpretable instructions with a virtual machine; and  
wherein said step of interpreting interprets said executing interpretable  
instructions, wherein said interrupted instruction is derived from said  
interpretable instructions, and  
said step of storing includes the step of executing the downloaded script to  
cause a user-mode trap to be sent to said interpretable program  
whereby virtual machine specific context information may be stored.
50. (Previously presented) The method of 16 further including the steps of:  
receiving a program having bytecode instructions;  
receiving a downloaded script that causes a user-mode trap to be sent to  
the program;  
executing the bytecode instructions with a virtual machine; and  
wherein said step of interpreting interprets said executing bytecode  
instructions, wherein said interrupted instruction is derived from said  
bytecode instructions, and  
said step of storing includes the step of executing the downloaded script to  
cause a user-mode trap to be sent to said program whereby virtual  
machine specific context information may be stored.

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

51. (Original) The method of claim 1 wherein said step of storing includes directly delivering the at least one data value of interest to the interrupted program.

52. (Original) The method of claim 51 wherein said step of storing includes generating a user-mode interrupt to deliver the at least one data value of interest to the interrupted program.

53. (Original) The method of claim 51 wherein a portion of the object code instructions are executed in a kernel, and said step of storing includes performing an upcall from the kernel to a user-mode handler in the interrupted program to store the at least one data value of interest.

54. (Original) The method of claim 1 wherein said step of storing includes the steps of:

- identifying a process identifier associated with the program; and
- storing the at least one data value of interest in a memory space associated with the process identifier such that only the program can access the at least one data value.

55. (Previously presented) The method of claim 1 wherein said step of storing includes the steps of:

- identifying access control identifiers associated with the program, particular ones of the access control identifiers also being associated a particular user; and
- storing the at least one data value such that only the particular user associated with the access control identifier can access the at least one data value.

56. (Original) The method of claim 55 wherein the access control identifier includes at least one of a process identifier, a user identifier and a group identifier.

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

57. (Original) The method of claim 1 wherein said step of storing includes the steps of:

identifying a group identifier associated with the program; and  
storing the at least one data value of interest in a user space associated with the group identifier such that only a user associated with the group identifier can access the at least one data value.

58. (Original) The method of claim 1 further comprising the steps of:  
storing the at least one data value of interest of the first database in a hotlist of most frequently occurring data values, the hotlist storing a predetermined number of data values and a count associated with each data value.

59. (Original) The method of claim 58 further comprising the step of encoding the hotlist.

60. (Original) The method of claim 59 wherein said step of encoding encodes the hotlist by sorting at least a subset of the data values of the hotlist by the count.

61. (Previously presented) The method of claim 59 wherein said step of encoding encodes the hotlist by reordering at least a subset of the data values of the hotlist such that the data values in the at least one subset are stored in contiguous memory locations.

62. (Original) The method of claim 58 wherein said step of storing the at least one data value of interest of the first database stores the at least one data value of interest using a randomized technique.

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

63. (Original) The method of claim 62 wherein said step of storing the at least one data value of interest dynamically adapts between a first randomized technique and a second randomized technique.

64. (Original) The method of claim 58 wherein said step of storing the at least one data value of interest of the first database in a hotlist of most frequently occurring data values uses a concise samples technique.

65. (Original) The method of claim 58 wherein said step of storing the at least one data value of interest of the first database in a hotlist of most frequently occurring data values uses a counting samples technique.

66. (Currently amended) A computer system for value sampling a computer program having object code instructions, comprising:

- a processor for executing the object code instructions of the computer program; and

- a memory for storing instructions that:

- deliver interrupts at time-based intervals during execution of the program, including delivering a first interrupt; and

- store at least one data value of interest in a first database in response to at least a subset of the first interrupts, the at least one data value of interest being associated with a particular object code instruction of the program, the particular object code instruction being executed by the computer when the interrupt occurs, and wherein the particular object code instruction of the program remains unmodified.

67. (Original) The computer system of claim 66 wherein said instructions that store further include instructions that:

- identify at least one issue block of instructions;

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

interpret the instructions of the at least one issue block; and  
store at least one data value of interest associated with at least one  
interpreted instruction.

68. (Original) The computer system of claim 66, wherein said instructions that deliver interrupts interrupt a particular object code instruction as an interrupted instruction, the particular object code instruction being associated with a memory address; and

wherein said instructions that store further include instructions that:  
store the memory address and the interrupted instruction;  
configure a second interrupt to be delivered after a predetermined  
number of instructions; and  
in response to the second interrupt,  
deactivate the second interrupt, and  
store the memory address and the at least one data value of  
interest for the associated instruction in the first  
database.

69. (Currently amended) A computer program product for value sampling a computer program having object code instructions, the computer program product for use in conjunction with a computer system, the computer program product comprising a computer readable storage medium and a computer program mechanism embedded therein, the computer program mechanism comprising:

Instructions that deliver interrupts at time-based intervals during execution of the program, including delivering a first interrupt;

instructions that, in response to at least a subset of the first interrupts, store at least one data value of interest in a first database, the at least one data value of interest being associated with a particular object code instruction of the program, the particular object code instruction being executed by the computer when the interrupt

**Appl. No. 09/401,616**  
**Amdt. dated January 21, 2005**  
**Reply to Office action of October 21, 2004**

occurs, and wherein the particular object code instruction of the program remains unmodified.

70. (Original) The computer program product of claim 69 wherein the intervals are random intervals.

71. (Original) The computer program product of claim 69 wherein the instructions that store further include instructions that:

- identify at least one issue block of instructions;
- interpret the instructions of the at least one issue block; and
- store at least one data value of interest associated with at least one interpreted instruction.

72. (Original) The computer program product of claim 69 wherein the instructions that deliver interrupts interrupt a particular object code instruction as an interrupted instruction, the particular object code instruction being associated with a memory address; and

wherein the instructions that store further include instructions that:

- store the memory address and the interrupted instruction;
- configure a second interrupt to be delivered after a predetermined number of instructions; and
- in response to the second interrupt,
  - deactivate the second interrupt, and
  - store the memory address and the at least one data value of interest for the associated instruction in the first database.

**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

**BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☒ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER: \_\_\_\_\_**

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.**